A framework for adaptive regularization in streaming Lasso models

Ricardo Pio Monti[1], Christoforos Anagnostopoulos[1], and Giovanni Montana[1,2]

[1]Department of Mathematics, Imperial College London, London, SW7 2AZ, UK
[2]Department of Biomedical Engineering, King's College London, St Thomas' Hospital, London, SE1 7EH, UK

arXiv:1610.09127v1 [stat.ML] 28 Oct 2016

Abstract

Large scale, streaming datasets are ubiquitous in modern machine learning. Streaming algorithms must be scalable, amenable to incremental training and robust to the presence of non-stationarity. In this work consider the problem of learning ℓ_1 regularized linear models in the context of streaming data. In particular, the focus of this work revolves around how to select the regularization parameter when data arrives sequentially and the underlying distribution is non-stationary (implying the choice of optimal regularization parameter is itself time-varying). We propose a framework through which to infer an adaptive regularization parameter. Our approach employs an ℓ_1 penalty constraint where the corresponding sparsity parameter is iteratively updated via stochastic gradient descent. This serves to reformulate the choice of regularization parameter in a principled framework for online learning and allows for the derivation of convergence guarantees in a non-stochastic setting. We validate our approach using simulated and real datasets and present an application to a neuroimaging dataset.

1 Introduction

We are interested in learning ℓ_1 regularized linear regression models in the context of streaming, non-stationary data. There has been significant research relating to the estimation of such models in a streaming data context [7, 8]. However a fundamental aspect that has been overlooked is the selection of the regularization parameter. The choice of this parameter dictates the severity of the regularization penalty. While the underlying optimization problem remains convex, distinct choices of such a parameter yield models with vastly different characteristics. This poses significant concerns from the perspective of model performance and interpretation. It therefore follows that selecting such a parameter is an important problem that must be addressed in a data-driven manner.

Many solutions have been proposed through which to select the regularization parameter in a non-streaming context. For example, stability based approaches have been proposed in the context of linear regression [25]. Other popular alternatives include cross-validation [11] and information theoretic techniques. However, in a streaming setting such approaches are infeasible due to the limited computational resources available. Moreover, the statistical properties of the data may vary over time; a common manifestation being concept drift [1, 36]. This complicates the use of sub-sampling methods as the data can no longer be assumed to follow a stationary distribution. Furthermore, as we argue in this work, it is conceivable that the optimal choice of regularization parameter may itself vary over time. It is also important to note that traditional approaches such change-point detection cannot be employed as there is no readily available pivotal quantity. It therefore follows that novel methodologies are required in order to tune regularization parameters in an online setting.

Applications involving streaming datasets are abundant, ranging from finance to cyber-security [18] and neuroscience [35]. In this work we are motivated by the latter application, where penalized regression models are often employed to decode statistical dependencies across spatially remote brain regions, referred to as functional connectivity [31]. A novel avenue for neuroscientific research involves the study of functional connectivity in real-time [35]. Such research faces challenges due to the non-stationary as well as potentially high dimensional nature of neuroimaging data [26]. In order to address these challenges, many of the proposed methods to date have employed fixed sparsity parameters [28]. However, such choices are typically justified only by the methodological constraints associated with updating the regularization parameter, as opposed to for biological reasons.

In order to address these issues we propose a framework through which to learn an adaptive sparsity parameter in an online fashion. The proposed framework, named Real-time Adaptive Penalization (RAP), is capable of iteratively learning time-varying regularization parameters via the use of adaptive filtering. Briefly, adaptive filtering methods are semi-parametric methods which employ information from recent observations to tune a parameter of interest. In this manner, adaptive filtering methods are capable of handling temporal variation which cannot easily be modeled explicitly [17]. The contributions of this work can be summarized as follows:

1. We propose and validate a framework through which to tune a time-varying sparsity parameter for ℓ_1 regularized linear models in real-time.
2. We provide theoretical convergence guarantees in a non-stochastic setting.
3. We provide empirical validation of the proposed framework using both simulated and real datasets and present an application to a neuroimaging dataset.

The remainder of this manuscript is organized as follows: related work is discussed in Section 2. We formally describe our problem in Section 3 and the proposed framework is introduced in Section 4. We provide empirical evidence based on real and simulated data in Section 5.

2 Related work

Regularized methods have established themselves as popular and effective tools through which to handle high-dimensional data [15]. Such methods employ regularization penalties as a mechanism through which to constraint the set of candidate solutions, often with the goal of enforcing specific properties such as parsimony. In particular, ℓ_1 regularization is widely employed as a convex approximation to the combinatorial problem of model selection. As a result of the convex nature of ℓ_1 penalties, efficient and highly scalable estimation algorithms can be derived [4].

However, the introduction of an ℓ_1 penalty requires the specification of the associated regularization parameter. The task of tuning such a parameter has primarily been studied in the context of non-streaming, stationary data. Stability selection procedures, introduced by [25], effectively look to by-pass the selection of a specific regularization parameter by instead fitting multiple models across sub-sampled data. Variables are subsequently selected according to the proportion of all models in which they are present. In this manner, stability selection is able to provide important theoretical guarantees while incurring an additional computational burden. Other popular approaches involve the use of cross-validation [10] or information theoretic techniques [32]. However, such methods have been reported to perform poorly in high-dimensional settings [20, 22, 34] and cannot easily be adapted to handle streaming data.

Online learning with the ℓ_1 constraints has also been studied extensively and many computationally efficient algorithms are available. A stochastic gradient descent algorithm is proposed by [7]. More generally, online learning of regularized objective functions has been studied extensively by [8] who propose a general class of computationally efficient methods based on proximal gradient descent. The aforementioned methods all constitute important advances in the study of sparse online learning algorithms. However, a fundamental issue that has been overlooked corresponds to the selection of the regularization parameters. As such, current methodologies are rooted on the assumption that the regularization parameter remains fixed. It follows that the regularization parameter may itself vary over time, yet selecting such a parameter in a principled manner is non-trivial. The focus of this work is to present and validate a framework through which to automatically select and update the regularization parameter in real-time. The framework presented in this work is therefore complementary and can be employed in conjunction with many of the preceding techniques.

More generally, the automatic selection of hyper-parameters has recently become an active topic in machine learning [30]. Interest in this topic has been catalyzed by the success of deep learning algorithms, which typically involve many such hyper-parameters. Sequential model based optimization (SMBO) methods such as Bayesian optimization employ a probabilistic surrogate to model the generalization performance of learning algorithms as samples from a Gaussian process [30], leading to expert level performance in many cases. It follows that such methods may be employed to tune regularization parameters in the context of penalized linear regression models. However, there are several important differences between the SMBO framework and the proposed framework. The most significant difference relates to the fact that the proposed framework employs gradient information in order to tune the regularization parameter while SMBO methods such as Bayesian optimization are rooted in the use of a probabilistic surrogate model. This allows the SMBO framework to be

applied in a wide range of settings while the proposed framework focuses exclusively on Lasso regression models. However, as we describe in this work, the use of gradient information makes the RAP framework ideally suited in the context of non-stationary, streaming data. This is in contrast to SMBO techniques, which typically assume the data is stationary.

3 Preliminaries

In this section we introduce the necessary ingredients to derive the proposed framework. We begin formally defining the problem addressed in this work in Section 3.1. Adaptive filtering methods are introduced in Section 3.2.

3.1 Problem set-up

In this work we are interested in streaming data problems. Here it is assumed that pairs (X_t, y_t) arrive sequentially over time, where $X_t \in \mathbb{R}^{p\times 1}$ corresponds to a p-dimensional vector of predictor variables and y_t is a univariate response. The objective of this work is to learn time-varying linear regression models from which to accurately predict future responses, y_{t+1}, from predictors, X_{t+1}. An ℓ_1 penalty, parameterized by $\lambda \in \mathbb{R}_+$, is introduced in order to encourage sparse solutions as well as to ensure the problem is well-posed from an optimization perspective. This corresponds to the Lasso model introduced by [32]. For a given choice of regularization parameter, λ, time-varying regression coefficients can be estimated by minimizing the following convex objective function:

$$L_t(\beta, \lambda) = \sum_{i=1}^{t} w_i \left(y_i - X_i^T \beta\right)^2 + \lambda ||\beta||_1, \tag{1}$$

where $w_i > 0$ are weights indicating the importance given to past observations [1, 16]. Typically, w_i decay monotonically in a manner which is proportional to the chronological proximity of the ith observation. For example, weights w_i may be tuned using a fixed forgetting factor or a sliding window.

In a non-stationary context, the optimal estimates of regression parameters, $\hat{\beta}_t$, may vary over time. The same argument can be posed in terms of the selected regularization parameter, λ. For example, this may arise due to changes in the underlying sparsity or changes in the signal-to-noise ratio. While there exists a wide range of methodologies through which to update regression coefficients in a streaming fashion, the choice of regularization parameter has been largely ignored. As such, the primary objective of this work is to propose a framework through which to learn time-varying regularization parameter in real-time. The proposed framework seeks to iteratively update the regularization parameter via stochastic gradient descent. As such, the proposed method is conceptually related to adaptive filtering theory [17], discussed below, which has been extensively employed in the context of streaming data [3].

3.2 Adaptive filtering

Filtering, as defined in [17], is the process through which information regarding a quantity of interest is assimilated using data measured up to and including time t. In many real-time applications, the quantity of interest is assumed to vary over time. The task of a filter therefore corresponds to effectively controlling the rate at which past information is discarded. Adaptive filtering methods provide an elegant method through which to handle a wide range of non-stationary behavior without having to explicitly model the dynamic properties of the data stream. Their popularity in recent years has grown, partly due to the fact that in many scenarios the latter approach is infeasible.

The simplest filtering methods discard information at a constant rate, for example determined by a fixed forgetting factor [13]. More sophisticated methods are able to exploit gradient information to determine the aforementioned rate. Such methods are said to be adaptive as the rate at which information is discarded varies over time. It follows that the benefits of adaptive methods are particularly notable in scenarios where the quantity of interest is highly non-stationary.

To further motivate discussion, we briefly review filtering in the context of fixed forgetting factors for streaming linear regression. In such a scenario, it suffices to store summary statistics for the mean and sample covariance. For a given fixed forgetting factor $r \in (0, 1]$, the sample mean can be recursively estimated as follows:

$$\bar{X}_t = \left(1 - \frac{1}{\omega_t}\right) \bar{X}_{t-1} + \frac{1}{\omega_t} X_t, \tag{2}$$

where ω_t is a normalizing constant defined as:

$$\omega_t = \sum_{i=1}^{t} r^{t-i} = r \cdot \omega_{t-1} + 1. \tag{3}$$

Similarly, the sample covariance can be learned iteratively:

$$S_t = \left(1 - \frac{1}{\omega_t}\right) S_{t-1} + \frac{1}{\omega_t}(X_t - \bar{X}_t)^T (X_t - \bar{X}_t). \tag{4}$$

It is clear that the value of r directly determines the adaptivity of a filter as well as its susceptibility to noise. However, in many practical scenarios the choice of r presents a challenge as it assumes some knowledge about the *degree* of non-stationarity of the system being modeled as well as an implicit assumption that this is constant [17]. Adaptive filtering methods address these issues by allowing r to be tuned online in a data-driven manner. This is achieved by quantifying the performance of current parameter estimates for new observations, X_{t+1}. Throughout this work we denote such a measure by $C(X_{t+1})$.

A favored approach is to define $C(X_{t+1})$ to be the residual error on unseen data [5]. Then assuming $\frac{\partial C(X_{t+1})}{\partial r}$ can be efficiently calculated, our parameter of interest can be updated in a stochastic gradient descent framework:

$$r_{t+1} = r_t - \epsilon \left. \frac{\partial C(X_{t+1})}{\partial r} \right|_{r=r_t} \tag{5}$$

Here ϵ is a small step-size parameter which determines the learning rate.

The objective of this work therefore corresponds to extending adaptive filtering methods to the domain of learning a time-varying regularization parameter for Lasso regression models.

4 Methods

As noted previously, the choice of parameter λ dictates the severity of the regularization penalty. Different choices of λ result in vastly different estimated models. While several data-driven approaches are available for selecting λ in an offline setting, such methods are typically not feasible for streaming data for two reasons. First, limited computational resources pose a practical restriction. Second, data streams are often non-stationary and rarely satisfy *iid* assumptions required for methods based on the bootstrap [1]. Moreover, it is important to note that traditional methods such as change point detection cannot be employed due to the absence of a readily available pivotal quantity for λ.

In this section we detail the proposed framework for real-time adaptive penalization (RAP) in the context of streaming Lasso models. We begin by outlining the RAP framework and deriving the necessary machinery in Section 4.1. Section 4.2 outlines the resulting algorithm and discusses computational considerations. We provide convergence guarantees in a non-stochastic setting in Section 4.3.

4.1 Proposed framework

We propose to learn a time-varying sparsity parameter in an adaptive filtering framework. This allows the proposed method to relegate the choice of sparsity parameter to the data. Moreover, by allowing λ_t to vary over time the proposed method is able to naturally accommodate datasets where the underlying sparsity may be non-stationary.

We define the empirical objective to be the look-ahead residual error, defined as:

$$C_{t+1} = C(X_{t+1}, y_{t+1}) = ||y_{t+1} - X_{t+1}\hat{\beta}_t(\lambda_t)||_2^2, \tag{6}$$

where we write $\hat{\beta}_t(\lambda_t)$ to emphasize the dependence of the estimated regression coefficients on the current value of the regularization parameter, λ_t. Following Section 3.2, the regularization parameter can be iteratively updated as follows:

$$\lambda_{t+1} = G(\lambda_t) = \lambda_t - \epsilon \frac{\partial C_{t+1}}{\partial \lambda_t}. \tag{7}$$

We note that for convinience we write $\frac{\partial C_{t+1}}{\partial \lambda_t}$ to denote the derivate of C_{t+1} with respect to λ evaluated at $\lambda = \lambda_t$ (i.e, $\frac{\partial C_{t+1}}{\partial \lambda}|_{\lambda=\lambda_t}$). We note that λ_t is bounded below by zero, in which case no regularization is applied, and above by $\lambda_t^{max} = \max_j \left\{ |\sum_{i=1}^{t} w_i y_i X_{i,j}| \right\}$, in which case all regression coefficients are set to zero [12].

The proposed framework requires only the specification of an initial sparsity parameter, λ_0, together with a stepsize parameter, ϵ. In this manner the proposed framework effectively replaces a fixed sparsity parameter with a stepsize parameter, ϵ. This is desirable as the choice of a fixed sparsity parameter is difficult to justify in the context of streaming, non-stationary data. Moreover, any choice of λ is bound to be problem specific. In comparison, we are able to interpret ϵ as a stepsize parameter in a stochastic gradient descent scheme. As a result, there are clear guidelines which can be followed when selecting ϵ [6].

Once the regularization parameter has been updated, estimates for the corresponding regression coefficients can be obtained by minimizing $L_{t+1}(\beta, \lambda_{t+1})$, for which there is a wide literature available [6, 7, 8]. The challenge in this work therefore corresponds to efficiently calculating the derivative in equation (7). Through the chain rule, this can be decomposed as:

$$\frac{\partial C_{t+1}}{\partial \lambda_t} = \frac{\partial C_{t+1}}{\partial \beta_t} \cdot \frac{\partial \beta_t}{\partial \lambda_t}. \tag{8}$$

The first term in equation (8) can be obtained by direct differentiation. In the case of the second term, we leverage the results of [9] and [29] who demonstrate that the Lasso solution path is piecewise linear as a function of λ. By implication, $\frac{\partial \beta_t}{\partial \lambda_t}$ must be piecewise constant. Furthermore, there is a simple, closed-form solution for $\frac{\partial \beta_t}{\partial \lambda_t}$.

Proposition 1. *[Adapted from [29]] In the context of ℓ_1 penalized linear regression models, the derivative $\frac{\partial \beta_t}{\partial \lambda_t}$ is piecewise constant and can be obtained in closed form.*

Proof. Since $\hat{\beta}_t$ minimizes the objective function specified in equation (1), it satisfies:

$$\nabla_\beta \left(L_t(\beta, \lambda) \right) |_{\beta=\hat{\beta}_t} \ni 0 \tag{9}$$

We must therefore have that:

$$\begin{aligned} \frac{\partial}{\partial \lambda} \left(\nabla_\beta L_t(\beta, \lambda)|_{\beta=\hat{\beta}_t} \right) &= 0 && (10) \\ &= \frac{\partial \beta}{\partial \lambda} \nabla_\beta^2 L_t(\beta, \lambda) + \text{sign}(\hat{\beta}_t) && (11) \\ &= \frac{\partial \beta}{\partial \lambda} \left(X_{1:t}^T W X_{1:t} \right) + \text{sign}(\hat{\beta}_t) && (12) \end{aligned}$$

Where W is a diagonal matrix with entries w_i. Rearranging equation (12) yields:

$$\begin{aligned} \frac{\partial \beta_t}{\partial \lambda} &= - \left(X_{1:t}^T W X_{1:t} \right)^{-1} \text{ sign}(\beta_t), && (13) \\ &= -(S_t)^{-1} \text{ sign}(\beta_t). && (14) \end{aligned}$$

□

From Proposition 1 we have that the derivative, $\frac{\partial C_{t+1}}{\partial \lambda_t}$, can be computed in closed form. Moreover, we note that the derivative in equation (13) is only defined over the active set of regression coefficients, $\mathcal{A}_t = \{i : (\hat{\beta}_t(\lambda_t))_i \neq 0\}$, and zero elsewhere. In practice we must therefore consider two scenarios:

- the active set is non-empty (i.e., $\mathcal{A}_t \neq \emptyset$). In this case equation (13) is well-defined.
- the active set is empty. In this case we proceed to take a step in the direction of the most correlated predictor: $\hat{j} = \underset{j}{\text{argmax}} \left\{ |\sum_{i=1}^{t} w_i y_i X_{i,j}| \right\}$. Thus we have that:

$$\left(\frac{\partial \beta_t}{\partial \lambda} \right)_i = \delta_{i,\hat{j}} \text{ sign} \left(\sum_{i=1}^{t} w_i y_i X_{i,\hat{j}} \right), \tag{15}$$

where $\delta_{i,j}$ is the dirac-delta function.

4.2 Streaming Lasso regression

At each iteration, a new pair (X_{t+1}, y_{t+1}) is received and employed to update both the time-varying regularization parameter, λ_t, as well as the corresponding estimate of regression coefficients, $\hat{\beta}_t(\lambda_t)$. The former involves computing the derivative $\frac{\partial C_{t+1}}{\partial \lambda_1}$ as outlined in Section 4.1. As noted in equation (14), a current estimate of the sample covariance matrix is sufficient. This may be recursively estimated using a fixed forgetting factor, as shown in equation (4). The latter involves solving a convex optimization problem which can be addressed in a variety of ways. In this work we look to iteratively estimate regression coefficients using coordinate descent methods [12]. Such methods are easily amenable to streaming data and allow us to exploit previous estimates as warm starts. In our experience, the use of warm starts leads to convergence within a handful of iterations. Psuedo-code detailing the proposed RAP framework is given in Algorithm 1.

Algorithm 1 **R**eal-time **A**daptive **P**enalization

Require: $\epsilon \in \mathbb{R}_+$ and $r \in (0, 1]$

for $t \leftarrow 1 \dots t, \dots$ **do**
  receive new (X_{t+1}, y_{t+1})
  if $\mathcal{A}_t \neq \emptyset$ **then**
    set $\frac{\partial \beta_t}{\partial \lambda_t}$ using equation (13)
  else
    set $\frac{\partial \beta_t}{\partial \lambda_t}$ using equation (15)
  set $\frac{\partial C_{t+1}}{\partial \lambda_t} = \frac{\partial C_{t+1}}{\partial \beta_t} \frac{\partial \beta_t}{\partial \lambda_t}$
  update $\lambda_{t+1} = \lambda_t - \epsilon \frac{\partial C_{t+1}}{\partial \lambda_t}$
  $\hat{\beta}_{t+1}(\lambda_{t+1}) = \underset{\beta}{\operatorname{argmin}} \{L_{t+1}(\beta, \lambda_{t+1})\}$

4.2.1 Computational considerations

With respect to the computational and memory demands, we note that the major expense incurred when calculating $\frac{\partial \beta_t}{\partial \lambda_t}$ involves inverting the sample covariance matrix. While only the dimensions corresponding to active variables need to be considered, this still corresponds to inverting a $|\mathcal{A}_t| \times |\mathcal{A}_t|$ matrix. It is possible to alleviate the computational burden by efficiently updating $(S_t)_{\mathcal{A}_t, \mathcal{A}_t}$ using the Sherman - Morrison formula. In this case, care must be taken to ensure that the support of $\mathcal{A}_t$ has not changed from iteration $t-1$ to t. If this is not the case (i.e., a regression coefficient has either added/removed from $\mathcal{A}_t$) then the inverse must be re-calculated.

However, computational and memory efficiency is paramount to streaming methods. The need to compute and store the inverse of the sample covariance is undesirable in the context of high-dimensional data. As a result, the following approximation is proposed:

$$\frac{\partial \beta_t}{\partial \lambda_t} \approx -\left(\operatorname{diag}\left(S_t\right)\right)^{-1} \operatorname{sign}\left(\beta_t\right). \tag{16}$$

Here a diagonal approximation to the sample covariance is employed, implying that only the diagonal elements of the sample covariance must be stored and inverted. Such approximations are frequently employed in streaming or large data applications [8]. This serves to reduce the computational burden of updating the sparsity parameter in the proposed manner. The approximate update therefore has a time and memory complexity that is proportional to the cardinality of the active set, $\mathcal{A}_t$.

4.3 Fixed point convergence

In this section we study the convergence properties of the proposed framework. We begin by noting that the update rule is piecewise non-expansive over the support of regularization parameter. We then show that iteratively applying equation (7) leads to convergence to a fixed point.

Recall that $G(\lambda_t) = \lambda_t - \epsilon \frac{\partial C_{t+1}}{\partial \lambda_t}$ is a self-mapping defined on the support $\Lambda = [0, \lambda_t^{max}]$. We study the behavior of iteratively applying the update rule $G(\lambda_t)$ for fixed new data pair (X_{t+1}, y_{t+1}). This corresponds to iteratively performing the gradient descent update to minimize residual error, C_{t+1}, for some fixed unseen pair, (X_{t+1}, y_{t+1}). While the proposed algorithm is stochastic in the sense that distinct random samples, (X_{t+1}, y_{t+1}), are employed at each update step, the results presented below provide reassuring guarantees in a non-stochastic setting. We note that such non-stochastic results are often presented when studying online algorithms.

Throughout the remainder of this section we abuse notation and write $\lambda_{t+1} = G(\lambda_t)$ to denote the result of applying the gradient update for t iterations. Our goal is to show that the limit of $\delta_t = |\lambda_{t+1} - \lambda_t|$ converges to zero as the number of iterations, t, increases.

First, we demonstrate that the support of the regularization parameter, Λ, can be partitioned into $p-1$ subsets where G is a contraction mapping. Then the mappings across subsets are studied to show that proposed algorithm does not exhibit periodic, expansive behavior. These two results are combined to obtain convergence to a fixed point.

Assumption 1. *The coefficient profiles for each regression coefficient, $\{\beta(\lambda) : \lambda \in \Lambda\}$, are monotone.*

Assumption 1 implies that $|\beta(\lambda_1)| < |\beta(\lambda_2)|$ for all $\lambda_1 > \lambda_2$. For any λ we define the corresponding active set as $\mathcal{A}(\lambda)$. Assumption 1 thus implies $\mathcal{A}(\lambda_1) \subset \mathcal{A}(\lambda_2)$ for all $\lambda_1 > \lambda_2$. We note that the assumption of monotonicity has been previously employed to study the properties of Lasso estimators [14]. Furthermore, this assumption can be verified in practice by checking that the inverse covariance is diagonally dominant [14].

Lemma 1. *Under Assumption 1, Λ can be divided into $p-1$ open subsets, $\{S_i\}_{i=1}^{p-1}$, where G is a contraction mapping.*

Proof. We assume without loss of generality that $\lambda_1 > \lambda_2$. We consider:

$$|G(\lambda_1) - G(\lambda_2)| = \left| \lambda_1 - \lambda_2 - \epsilon \left(\frac{\partial C_{t+1}}{\partial \lambda_1} - \frac{\partial C_{t+1}}{\partial \lambda_2} \right) \right|. \tag{17}$$

Our objective is to show that $\frac{\partial C_{t+1}}{\partial \lambda_1} - \frac{\partial C_{t+1}}{\partial \lambda_2} > 0$, thereby showing that G is a contraction for suitably chosen ϵ. Recall the gradient with respect to regularization parameter λ is defined as:

$$\frac{\partial C_{t+1}}{\partial \lambda} = (y_{t+1} - X_{t+1}\beta_t(\lambda))^T X_{t+1}^T (S_t)^{-1} \operatorname{sign}(\beta_t(\lambda))$$

Due to Assumption 1, we have that:

$$\begin{aligned} \frac{\partial C_{t+1}}{\partial \lambda_1} - \frac{\partial C_{t+1}}{\partial \lambda_2} = & \underbrace{\sum_{i \in \mathcal{A}(\lambda_1) \cap \mathcal{A}(\lambda_2)} \left[(\beta_t(\lambda_2) - \beta_t(\lambda_1))^T \left(X_{t+1}^T X_{t+1} \right) (S_t)^{-1} \operatorname{sign}(\beta_t(\lambda_1)) \right]_i}_{A_1} \\ & \underbrace{- \sum_{i \in \mathcal{A}(\lambda_2) \setminus \mathcal{A}(\lambda_1)} \left[(y_{t+1} - X_{t+1}\beta_t(\lambda_2))^T X_{t+1}^T (S_t)^{-1} \operatorname{sign}(\beta_t(\lambda_2)) \right]_i}_{A_2} \end{aligned}$$

We note that A_2 will be zero whenever $\mathcal{A}(\lambda_1) \setminus \mathcal{A}(\lambda_2) = \emptyset$. Moreover, the term A_1 will always be greater than or equal to zero. This follows from the fact that $A_1 = g(\lambda_1) - g(\lambda_2)$ where

$$\begin{aligned} g(\lambda) &= - \left(\beta_t(\lambda)^T \left(X_{t+1}^T X_{t+1} \right) (S_t)^{-1} \operatorname{sign}(\beta_t(\lambda)) \right) \\ &= \beta_t(\lambda)^T \left(X_{t+1}^T X_{t+1} \right) \frac{\partial \beta_t(\lambda)}{\partial \lambda}. \end{aligned}$$

Therefore, we have that:

$$\frac{\partial g(\lambda)}{\partial \lambda} = \left(\frac{\partial \beta_t(\lambda)}{\partial \lambda}^T \left(X_{t+1}^T X_{t+1} \right) \frac{\partial \beta_t(\lambda)}{\partial \lambda} \right) \geq 0, \tag{18}$$

due to the positive semi-definite nature of $X_{t+1}^T X_{t+1}$. This indicates that $g(\lambda)$ is a monotone, non-decreasing function in λ. As a result, we are able to divide Λ into $p-1$ open subsets where the update rule G is a contraction for suitably selected ϵ. These subsets correspond to the regions where the support of the Lasso solution is constant, thus implying that A_2 is zero.

□

By Lemma 1, we have that $|G(\lambda_1) - G(\lambda_2)| < |\lambda_1 - \lambda_2|$ for all $\lambda_1, \lambda_2 \in S_i$. It remains to study the (possibly expansive) behavior across the subsets $\{S_i\}_{i=1}^{p-1}$; in particular there remains a need to mitigate against potential periodic, expansive behavior.

Lemma 2. *If periodic behavior occurs across subsets, then this must be a contraction.*

Proof. We consider periodic behavior of the form:

$$G(\lambda_t) \in \begin{cases} S_j & \text{if } t \text{ is even} \\ S_{j-1} & \text{if } t \text{ is odd} \end{cases} \tag{19}$$

We consider two subsets which we label S_1 and S_2. Without loss of generality we assume that $S_1 > S_2$ in the sense that $\lambda_1 > \lambda_2$ for all $\lambda_1 \in S_1$ and $\lambda_2 \in S_2$. We consider the periodic behavior described in equation (19).

Therefore, at an odd iteration the gradient update maps from S_2 into S_1. Thus we have $\lambda_t = G(\lambda_{t-1}) > \lambda_{t-1}$ by construction. This implies that $\frac{\partial C_t}{\partial \lambda_{t-1}} < 0$. Conversely, in every even iteration the gradient update maps from S_1 into S_2, implying that $\lambda_t = G(\lambda_{t-1}) < \lambda_{t-1}$. This in turn implies that $\frac{\partial C_t}{\partial \lambda_{t-1}} > 0$.

As a result, we have that for any $\lambda_1 \in S_1$ and $\lambda_2 \in S_2$:

$$\frac{\partial C_t}{\partial \lambda_1} - \frac{\partial C_t}{\partial \lambda_2} > 0$$

indicating that cyclic mapping must be contractions. □

By Lemma 2 we have that any periodic behavior across subsets must be a contraction. Due to the compact nature of Λ, it follows that at most $p-1$ (possibly expansive) non-periodic mappings across subsets occur, after which only contraction mappings occur.

Proposition 2. *Iteratively applying the gradient descent mapping G over a fixed training example, (X_{t+1}, y_{t+1}), leads to convergence to a fixed point.*

Proof. We consider the sequence $\{\delta_t\}_t$ where $\delta_t = |\lambda_{t+1} - \lambda_t|$. The terms in $\{\delta_t\}_t$ can be split exactly into two subsequences; one containing all mappings within the same subset S_i for some i and another containing all mappings across subsets. We denote these subsequences by $\{\delta_t\}_{t_{(1)}}$ and $\{\delta_t\}_{t_{(2)}}$ respectively.

By Lemma 1, the first subsequence consists of purely contraction mappings and therefore converges to zero. Similarly, Lemma 2 states that all periodic behavior across subsets must be a contraction, thereby implying that the second subsequence also converges to zero. As both subsequences contain all elements of $\{\delta_t\}_t$ and converge to zero, it follows that $\{\delta_t\}_t$ also converges to zero. □

We note that the aforementioned results also hold when either the exact or approximate gradient as well as when multiple unseen samples $\{(X_i, y_i) : i = 1, \ldots, T\}$ are employed (in the case of mini-batch updates).

Software

All algorithms described above have been implemented in an R package, `rRAP`, and can be downloaded together with documentation from the Comprehensive R Archive Network (CRAN).

5 Empirical results

In this section we look to empirically demonstrate the capabilities of the proposed framework by studying a variety of real and simulated datasets. In order to provide a flavor for the capabilities of the RAP algorithm, we being by studying the `diabetes` dataset in Section 5.1. This corresponds to a publicly available dataset which has been widely studied in the context of Lasso models [9]. We complement these results with a more extensive simulation study presented in Section 5.2.

5.1 Diabetes dataset

We begin by considering the `diabetes` dataset presented in [9]. Here the response is a quantitative measure of disease progression one year after baseline. The covariates associated with each observation correspond to measurements over ten baseline variables.

The objective of this example is to provide empirical evidence that the RAP algorithm is able to reliably track the regularization parameter. As such, cross-validation with $K = 10$ folds was employed in order to estimate the regularization parameter. The regularization parameter was subsequently estimated using the RAP framework in a streaming fashion. The RAP algorithm was

applied over $N = 500$ iterations. At each iteration, the dataset rows were randomly permuted such that the order in which observations arrived varied. Due to the stationary nature of the data, dynamically tracking regression coefficients was not of interest here. As such, the sample covariance was recursively estimated using a fixed forgetting factor of $r = 1$. This corresponds to an online analysis where information from past observations is not discarded. For each iteration, the proposed method was initialized with $\beta = \mathbf{0} \in \mathbb{R}^{p \times 1}$ Both the exact and approximate gradient updates were considered. Finally, a SMBO approach, in the form of Bayesian optimization, was also considered. This involved modeling the generalization performance of the penalized regression model as a Gaussian process with a square exponential covariance function. The expected improvement acquisition function was employed to search the parameter space for λ [30].

The results are shown in Figure [1], where the ℓ_1 norm[1] is plotted against the number of iterations of the RAP algorithm. The horizontal red and blue lines indicates the regularization parameter as selected by 10-fold cross-validation and SMBO respectively. The dashed lines represent the mean ℓ_1 norm selected over $N = 500$ permutations when the exact (black) or approximate (brown) gradients where employed. We note that in both cases the estimated ℓ_1 norm quickly increases away from zero and converges to the cross-validated norm.

[Figure 1 about here.]

5.2 Simulation study

In this section we look to compliment the results presented in Section 5.1 with a more extensive set of simulations. We begin by considering the performance of the RAP algorithm in the context of stationary data. This simulation serves to demonstrate that the proposed method is capable of accurately tracking the regularization parameter. We then study the performance of RAP algorithm in the context of non-stationary data. Throughout this simulation study the RAP algorithm is benchmarked against two offline methodologies: cross-validation and SMBO. In the context of SMBO methods, we study the performance against Bayesian optimization methods. Here a Gaussian process with a square exponential kernel was employed as a surrogate model together with the expected improvement acquisition function.

We begin by outlining simulation settings as well as the associated performance metrics in Sections 5.2.1 and 5.2.2 respectively. The results for the stationary and non-stationary data simulations are subsequently presented in Section 5.2.3 and 5.2.4.

5.2.1 Simulation settings

In order to thoroughly test the performance of the RAP algorithm, we look to generate synthetic data were we are able to control both the underlying structure as well as the dimensionality of the data. In this work, data was generated according to a multivariate Gaussian distribution with a block covariance structure. This introduced significant correlations across covariates, thereby increasing the difficultly of the regression task. Formally, the data simulation process followed that described by [23]. This involved sampling each covariate as follows:

$$X_t \sim \mathcal{N}(0, \Sigma),$$

where $\Sigma \in \mathbb{R}^{p \times p}$ is a block diagonal matrix consisting of five equally sized blocks. Within each block, the off-diagonal entries were fixed at 0.8, while the diagonal entries were fixed to be one. Having generated covariates, X_t, a sparse vector of regression coefficients, β, was simulated. This involved randomly selecting a proportion, ρ, of coefficients and randomly generating their values according to a standard Gaussian distribution. All remaining coefficients were set to zero. Finally, univariate responses were obtained as follows:

$$y_t \sim \mathcal{N}(X_t\beta, 1).$$

In this manner, it is possible to generate piece-wise stationary data, $\{(y_t, X_t) : t = 1, \ldots, T\}$. When studying the performance of the RAP algorithm in the context of stationary data, it sufficed to simulate one such dataset. In order to quantify performance in the context of non-stationary data, we concatenate multiple piece-wise stationary datasets. This results in datasets with abrupt changes. We note that in the non-stationary setting the block structure was randomly permuted at each iteration in order to avoid covariates constantly sharing the same set of highly correlated variables.

[1]The ℓ_1 norm was considered as opposed to the estimated sparsity parameter in order to avoid potential confusion arising due to scaling of regularization parameters and other idiosyncrasies. There is a one-to-one relationship between the sparsity parameter, λ, and the ℓ_1 norm.

5.2.2 Performance metrics

In order to asses the performance of the RAP algorithm we consider various performance metrics. In the context of stationary data, our primary objective is to demonstrate that the proposed method is capable of tracking the regularization parameter when benchmarked against traditional methods such as cross-validation. As a result, we consider the difference in ℓ_1 norms of the regression model estimated by each algorithm. This is defined as:

$$\Delta = ||\beta(\lambda^{CV})||_1 - ||\beta(\lambda^{RAP})||_1, \tag{20}$$

where we write λ^{CV} and λ^{RAP} to denote the regularization parameters selected by cross-validation and RAP algorithms respectively. We choose to employ the ℓ_1 norm (as opposed to directly considering the sparsity parameter, λ) as there is a one-to-one relationship between λ and the ℓ_1 norm. This serves to bypass any potential issues arising from scaling or other idiosyncrasies.

In the context of non-stationary data we are interested in two additional metrics. The first corresponds to the residual error over unseen observations, C_{t+1}, initially defined in equation (6). Secondly, we also consider the correct recovery of the sparse support of β_t. In this context, we treat the recovery of the support of β_t as a binary classification problem and quantify the performance using the F score; defined as the harmonic mean between the precision and recall of a classification algorithm.

5.2.3 Stationary data

While the results presented in Section 5.1 provide reassuring empirical evidence, we consider a more extensive simulation study here. In particular, we study the performance of the RAP algorithm as the dimensionality of regression coefficients, p, increases. The goal of this simulation therefore is to demonstrate that proposed algorithm is able to accurately track the regularization parameter.

Data was generated as described in Section 5.2.1 and the dimensionality of the covariates, X_t, was varied from $p = 10$ through to $p = 100$. For each value of p, $N = 500$ datasets were randomly generated. The regularization parameter was first estimated using $K = 10$ fold cross-validation. The RAP algorithm was subsequently employed and the difference in ℓ_1, defined in equation (20), was then computed.

This procedure was repeated to produce data with dimensionality $p = 10$ through to $p = 100$. For each simulated dataset, the regularization parameter was selected via $K = 10$ fold cross-validation in an offline manner (i.e., by considering the entire dataset). The RAP algorithm was subsequently employed to estimate the regularization parameter. This involved iterating through observations in a streaming fashion.

The difference in selected regularization parameters over $N = 500$ simulations is visualized in Figure [2]. It is reassuring to note that the differences are both small in magnitude as well as centered around the origin. The latter serves to indicate the absence of a large systematic bias. The figure also does not show evidence of any systematic change in the bias as the dimensionality increases.

[Figure 2 about here.]

5.2.4 Non-stationary data

While Section 5.2.3 provided empirical evidence demonstrating that the RAP framework can be effectively employed to track regularization parameters in a stationary setting, we are ultimately interested in streaming, non-stationary datasets. As a result, in this simulation we study the performance of the proposed framework in the context of non-stationary data.

While there are a multitude of methods through which to simulate non-stationary data, in this simulation study we chose to generate data with piece-wise stationary covariance structure. As a result, the underlying covariance alternated between two regimes: a sparse regime where the response was driven by a reduced subset of covariates and a dense regime where the converse was true. Thus, pairs (y_t, X_t) of response and predictors were simulated in a piece-wise stationary regimes. The dimensionality of the covariates was fixed at $p = 20$, implying that $X_t \in \mathbb{R}^{20}$. Changes occurred abruptly every 100 observations and two change-points were considered, resulting in 300 observations in total.

Covariates, X_t, were simulated as described in Section 5.2.1 within two alternating regimes; dense and sparse. The block-covariance structure remained fix within each regime (i.e., for 100 observations). Within the dense regime, a proportion $\rho_1 = 0.8$ of regression coefficients were randomly selected and their values sampled from a standard Gaussian distribution. All remaining coefficients

were set to zero. Similarly, in the case of the sparse regime, $\rho_2 = 0.2$ regression coefficients were randomly selected with remaining coefficients set to zero. The regression coefficients remained fixed within each regime.

In order to benchmark the performance of the proposed RAP framework, streaming penalized Lasso models were also estimated using a fixed and piece-wise constant sparsity parameters. As a result, the RAP algorithm was benchmarked against three distinct offline methods for selecting the regularization parameter. In the case of a fixed sparsity parameter, $K = 10$ fold cross-validation as well as Bayesian optimization were employed. Finally, cross-validation was also employed to learn a piece-wise constant regularization parameter. This was achieved by performing cross-validation for the data within each regime. For each of these methods, their offline nature dictated that the entire dataset should be analyzed simultaneously (as opposed to in a streaming fashion by the RAP algorithm). As such, they serve to provide a benchmark but would infeasible in the context of streaming data.

Results are shown in Figure [3], where the left panel shows the estimated regularization parameter as a function of time. These results provide further evidence to indicate that the RAP algorithm is able to reliably track the piece-wise constant parameters selected by cross-validation. As expected, there is some lag directly after each change occurs, however, the estimated regression parameters quickly adapt. The right panel of Figure [3] shows the mean residual error, C_{t+1}, for unseen data. We note there are abrupt spikes every 100 observations, corresponding to the abrupt changes in the underlying dependence structure. Detailed results are provided in Table 1. We note that the proposed framework is able to outperform the alternative offline approaches. In the case of the offline cross-validation and SMBO, this is to be expected as a fixed choice of regularization parameter is misspecified.

[Figure 3 about here.]

[Table 1 about here.]

6 Application to fMRI data

In this section we present an application of the RAP algorithm to task-based functional MRI (fMRI) data. This data corresponds to time-series measurements of blood oxygenation, a proxy for neuronal activity, taken across a set of spatially remote brain regions [21]. Our objective in this work is to quantify pairwise statistical dependencies across brain regions, typically referred to as functional connectivity within the neuroimaging literature [31].

While traditional analysis of functional connectivity was rooted on the assumption of stationarity, there is growing evidence to suggest this is not the case [19]. This particularly true in the context of task-based fMRI studies. Several methodologies have been proposed to address the non-stationary nature of fMRI data [2, 27, 26], many of which are premised on the use of penalized regression models such as those studied in this work. While such methods have made important progress in the study of non-stationary connectivity networks, they have typically employed fixed regularization parameters. This is difficult to justify in the context of non-stationary data and plausible biological justifications are not readily available. The RAP algorithm is therefore ideally suited to both accurately estimating non-stationary connectivity structure as well as providing insight regarding whether the assumption of a fixed sparsity parameter is reasonable.

6.1 Estimating connectivity via Lasso regressions

Estimating functional connectivity networks is fundamentally a statistical challenge. A functional relationship is said to exist across two spatially remote brain regions if their corresponding time-series share some statistical dependence. While this can be quantified in a variety of ways, a popular approach is the use of Lasso regression models to infer the conditional independence structure of a particular node. In such an approach, the time-series of a given node is regressed against the time-series of all remaining nodes. A functional relationship is subsequently inferred between the target node and all remaining nodes associated with a non-zero regression coefficient. The connectivity structure across all nodes can then be inferred via a neighbourhood selection approach [24]. The proposed RAP framework can directly be incorporated into such a model, resulting in time-varying conditional dependence structure where the underlying sparsity parameter is also inferred.

6.2 HCP Emotion Task Data

Emotion task data from the Human Connectome Project (HCP) [33] was studied with 20 subjects selected. During the task participants were presented with blocks of trials that either required them to decide which of two faces presented on the bottom of the screen match the face at the top of the screen, or which of two shapes presented at the bottom of the screen match the shape at the top of the screen. The faces had either an angry or fearful expression while the shapes represented the emotionally neutral condition. Preprocessing involved regression of Friston's 24 motion parameters from the fMRI data. Sixty-eight cortical and 16 subcortical ROIs were derived from the Desikan-Killiany atlas and the ASEG atlas, respectively. Mean BOLD time series for each of these 84 ROIs were extracted and further cleaned by regressing out time series sampled from white matter and cerebrospinal fluid. Finally, the extracted time courses were high-pass filtering using a cut-off frequency of $\frac{1}{130}$ Hz. Neurosynth, a platform for large-scale automated synthesis of neuroimaging data,was employed to reduce the number of regions studied. This provided an automatically generated forward inference map based on 790 studies quantifying the activation all regions in emotion studies. Twenty regions identified as core emotion hubs were selected. Data for each subject therefore consisted of $n = 175$ observations across $p = 20$ nodes.

6.3 Results

Data for each subject was analyzed independently where the time-varying estimates of the conditional dependence structure for each node were estimated as described in Section 6.1. A fixed forgetting factor of $r = .95$ was employed throughout with a stepsize parameter $\eta = .025$. The exact gradient was employed when updating the sparsity parameter at each iteration. In order to avoid unreliable initial performance of the algorithms a burn-in of twenty observations was employed.

The mean sparsity parameter over all subjects is shown in the top panel Figure 4. We observe decreased sparsity parameters for blocks in which subjects were presented with emotional (i.e., angry or fearful) faces (top panel, purple shaded areas) as compared to blocks in which subjects were shown neutral shapes (top panel, green shaded areas). The oscillation in sparsity parameter is highly correlated with task onset. When inspecting the networks estimated using the time varying sparsity parameter (bottom panel), we find strong coupling amongst many of the regions during the emotion processing blocks (A and C) compared to a clearly sparser network representation for blocks that require no emotion processing (i.e., neutral shapes, block B). This is to be expected as the selected regions are core hubs involved with emotion processing; therefore explaining the higher network activity during the emotion task when compared to the neutral task.

[Figure 4 about here.]

7 Conclusion

In this work we have presented a framework through which to learn a time-varying sparsity parameters in the context of streaming Lasso models. An approximate algorithm is also provided to address issues concerning computational efficiency; a factor of paramount importance in the context of high-dimensional data. We provide theoretical results regarding the convergence in a non-stochastic scenario. These results hold for both the exact and approximate gradient algorithms as well as in the context of mini-batch updates. Finally, empirical evidence is provided to validate the proposed algorithm.

We present two simulation studies which demonstrate the capabilities of the proposed method. These simulations demonstrate that the proposed RAP framework is capable of tracking the regularization parameter both in a stationary as well as non-stationary context. Finally, we present an application to task-based fMRI data, which is widely accepted to be non-stationary [19].

Future work will involve extending the RAP framework to consider alternative regularization schemes. In particular, the popular ridge or ℓ_2 penalty could be incorporated as the derivative, $\frac{\partial \beta}{\partial \lambda}$, is also available in closed form. Furthermore, it would also be possible to extend the framework to consider a wider range of models. In this setting, the results of [29] could be leveraged to consider time-varying sparsity for generalized linear models and support vector machines.

Acknowledgements

The authors wish to thank Romy Lorenz and Rob Leech for help in processing the data as well as helpful discussions.

References

[1] C. Aggarwal. *Data streams: models and algorithms*, volume 31. Springer Science & Business Media, 2007.

[2] E. Allen, E. Damaraju, S. Plis, E. Erhardt, T. Eichele, and V. Calhoun. Tracking whole-brain connectivity dynamics in the resting state. *Cerebral cortex*, page bhs352, 2012.

[3] C. Anagnostopoulos, D. Tasoulis, N. Adams, N. Pavlidis, and D. Hand. Online linear and quadratic discriminant analysis with adaptive forgetting for streaming classification. *Statistical Analysis and Data Mining*, 5(2):139–166, 2012.

[4] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. Optimization with sparsity-inducing penalties. *Foundations and Trends® in Machine Learning*, 4(1):1–106, 2012.

[5] A. Benveniste, M. Métivier, and P. Priouret. *Adaptive algorithms and stochastic approximations*, volume 22. Springer Science & Business Media, 2012.

[6] L. Bottou. Online learning and stochastic approximations. *On-line learning in neural networks*, 17(9):142, 1998.

[7] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *COMPSTAT'2010*, pages 177–186. Springer, 2010.

[8] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.

[9] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.

[10] J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani. Pathwise coordinate optimization. *The Annals of Applied Statistics*, 1(2):302–332, 2007.

[11] J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.

[12] J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1, 2010.

[13] F. Gustafsson. *Adaptive filtering and change detection*, volume 1. Wiley New York, 2000.

[14] T. Hastie, J. Taylor, R. Tibshirani, and G. Walther. Forward stagewise regression and the monotone lasso. *Electronic Journal of Statistics*, 1:1–29, 2007.

[15] T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical learning with sparsity: the lasso and generalizations*. CRC Press, 2015.

[16] M. Hayes. *Statistical digital signal processing and modeling*. John Wiley & Sons, 2009.

[17] S. Haykin. *Adaptive filter theory*. Pearson, 2008.

[18] N. Heard, D. Weston, K. Platanioti, and D. Hand. Bayesian anomaly detection methods for social networks. *The Annals of Applied Statistics*, 4(2):645–662, 2010.

[19] M. Hutchison, T. Womelsdorf, E. Allen, P. Bandettini, V. Calhoun, M. Corbetta, S. Della Penna, Jeff H Duyn, Gary H Glover, Javier Gonzalez-Castillo, et al. Dynamic functional connectivity: promise, issues, and interpretations. *NeuroImage*, 80:360–378, 2013.

[20] C. Leng, Y. Lin, and G. Wahba. A note on the lasso and related procedures in model selection. *Statistica Sinica*, pages 1273–1284, 2006.

[21] M. Lindquist. The statistical analysis of fmri data. *Statistical Science*, 23(4):439–464, 2008.

[22] H. Liu, K. Roeder, and L. Wasserman. Stability approach to regularization selection (stars) for high dimensional graphical models. In *NIPS*, pages 1432–1440, 2010.

[23] B. McWilliams, C. Heinze, N. Meinshausen, G. Krummenacher, and H. Vanchinathan. LOCO: Distributing ridge regression with random projections. *arXiv preprint arXiv:1406.3469*, 2014.

[24] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*, pages 1436–1462, 2006.

[25] N. Meinshausen and P. Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B*, 72(4):417–473, 2010.

[26] R. P. Monti, P. Hellyer, D. Sharp, R. Leech, C. Anagnostopoulos, and G. Montana. Estimating time-varying brain connectivity networks from functional MRI time series. *NeuroImage*, 103: 427–443, 2014.

[27] R. P. Monti, R. Lorenz, P. Hellyer, R. Leech, C. Anagnostopoulos, and G. Montana. Graph embeddings of dynamic functional connectivity reveal discriminative patterns of task engagement in HCP data. *Pattern Recognition in NeuroImaging (PRNI), 2015 International Workshop on*, pages 1–4, 2015. doi: 10.1109/PRNI.2015.21.

[28] R. P. Monti, R. Lorenz, R. Braga, C. Anagnostopoulos, R. Leech, and G. Montana. Real-time estimation of dynamic functional connectivity networks. *Human Brain Mapping*, 2016.

[29] S. Rosset and J. Zhu. Piecewise linear regularized solution paths. *The Annals of Statistics*, pages 1012–1030, 2007.

[30] B. Shahriari, K. Swersky, Z. Wang, R. Adams, and N. de Freitas. Taking the human out of the loop: A review of bayesian optimization. *IEEE Proceedings*, 104(1):148–175, 2016.

[31] S. Smith, K. Miller, G. Salimi-Khorshidi, M. Webster, C. Beckmann, T. Nichols, J. Ramsey, and M. Woolrich. Network modelling methods for fMRI. *NeuroImage*, 54(2):875–891, 2011.

[32] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B*, pages 267–288, 1996.

[33] D. Van Essen, S. Smith, D. Barch, T. Behrens, E. Yacoub, K. Ugurbil, and WU-Minn HCP Consortium. The WU-Minn human connectome project: an overview. *NeuroImage*, 80:62–79, 2013.

[34] L. Wasserman and K. Roeder. High dimensional variable selection. *Annals of statistics*, 37(5A): 2178, 2009.

[35] N. Weiskopf. Real-time fMRI and its application to neurofeedback. *NeuroImage*, 62(2):682–692, 2012.

[36] G. Widmer and M. Kubat. Learning in the presence of concept drift and hidden contexts. *Machine learning*, 23(1):69–101, 1996.

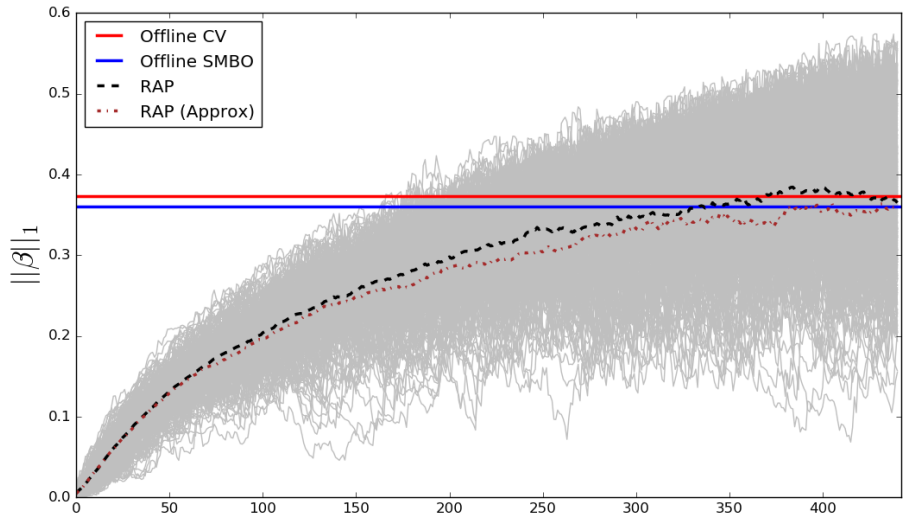

Figure 1: Mean ℓ_1 norm over all permutations for exact and approximate RAP algorithms is plotted in the dashed black and brown lines respectively. Each grey line corresponds to the ℓ_1 norm over each of the $N = 500$ iterations when applying RAP algorithm with exact gradient updates. The solid, horizontal lines correspond to the regularization parameters selected via two offline methods: cross-validation and Bayesian optimization.

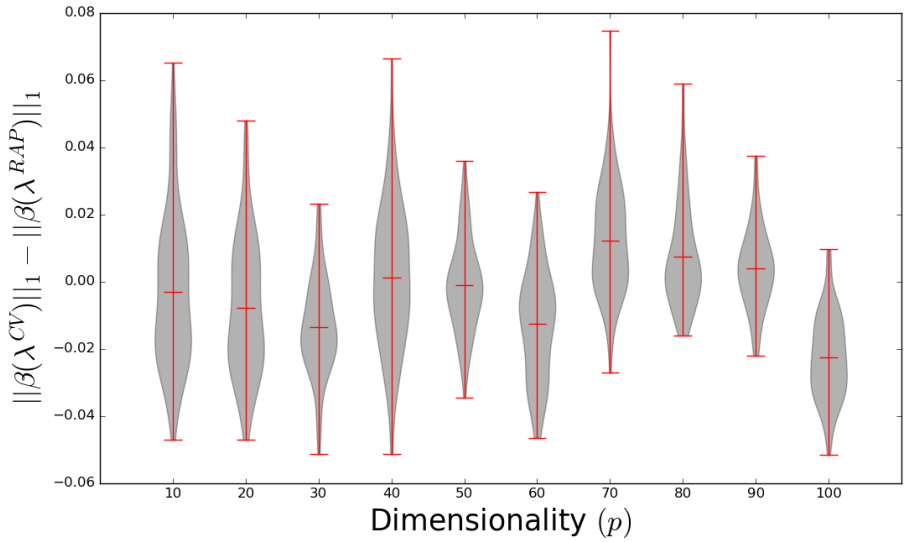

Figure 2: Violin plots visualizing the difference in selected regularization parameters as a function of the dimensionality, p. It is reassuring to note that the difference is both small in magnitude and centered around the origin, indicating the absence of a large systematic bias.

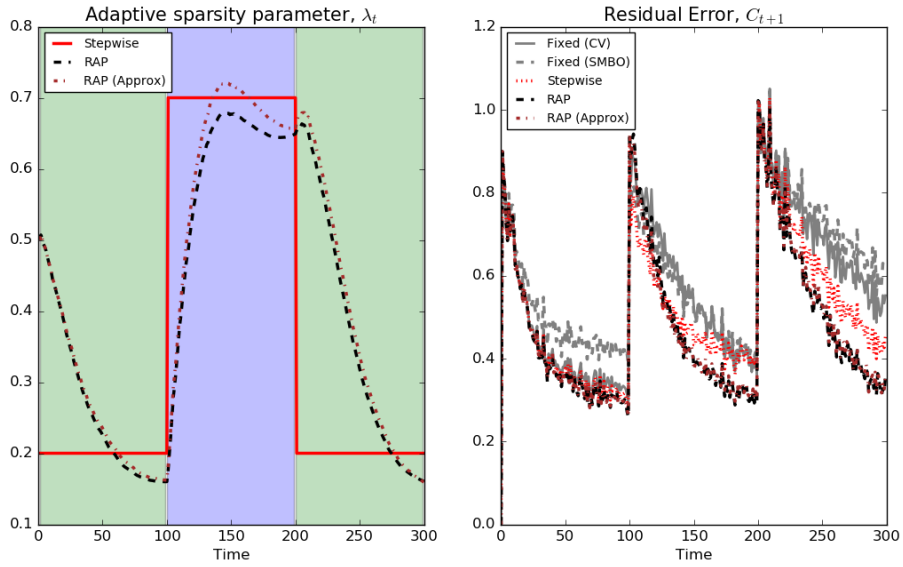

Figure 3: *Left*: Mean estimates of the regularization parameter are shown for the RAP algorithm as well as a the optimal piece-wise constant value selected by cross-validation. The background color indicates the nature of the underlying regime (green indicating dense, blue sparse). *Right*: Residual error, C_{t+1}, is plotted as a function of time. We note that the RAP algorithms outperforms the offline approaches employed.

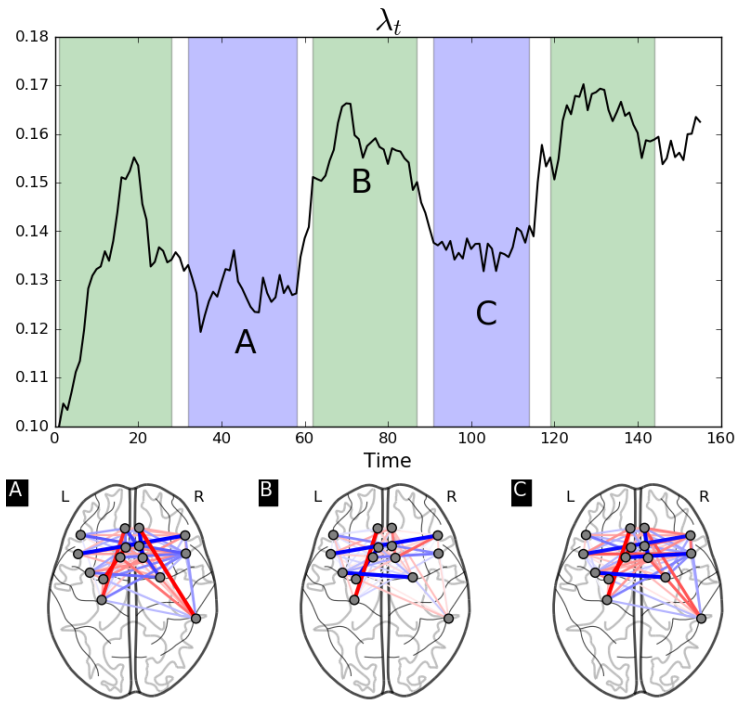

Figure 4: *Top*: the mean sparsity parameter is shown as a function of time. The background color indicates the nature of the task at hand (green indicates neutral task while blue indicates the emotion task). *Bottom*: estimated networks visualizing the estimated connectivity structure at three distinct points in time. Edge colors indicate the nature of the dependence (blue indicates a positive dependence, red a negative dependence).

Algorithm	$\bar{C}_t$	$\bar{F}_t$
Fixed (CV)	0.58 (0.05)	0.49 (0.05)
Fixed (SMBO)	0.63 (0.05)	0.50 (0.07)
Piecewise	0.51 (0.04)	0.56 (0.04)
RAP	**0.47** (0.04)	**0.64** (0.06)
RAP (Approx)	**0.48** (0.05)	**0.63** (0.07)

Table 1: Detailed results consisting of the mean negative log-likelihood, $\bar{C}_t$, as well as the mean F-score, $\bar{F}_t$. Standard errors are provided in brackets.